

POST-FLIGHT EDL ENTRY GUIDANCE PERFORMANCE OF THE 2011 MARS SCIENCE LABORATORY MISSION

Gavin F. Mendeck* and Lynn Craig McGrew†

The 2011 Mars Science Laboratory was the first Mars guided entry which safely delivered the rover to a landing within a touchdown ellipse of 19.1 km x 6.9 km. The Entry Terminal Point Controller guidance algorithm is derived from the final phase Apollo Command Module guidance and, like Apollo, modulates the bank angle to control the range flown. The guided entry performed as designed without any significant exceptions. The Curiosity rover was delivered about 2.2 km from the expected touchdown. This miss distance is attributed to little time to correct the downrange drift from the final bank reversal and a suspected tailwind during heading alignment. The successful guided entry for the Mars Science Laboratory lays the foundation for future Mars missions to improve upon.

INTRODUCTION

The 2011 Mars Science Laboratory (MSL) was the first Mars guided entry which safely delivered the rover to a final position approximately 2 km from its target within a touchdown ellipse of 19.1 km x 6.9 km. The Entry Terminal Point Controller guidance algorithm is derived from the final phase Apollo Command Module guidance and, like Apollo, modulates the bank angle to control the range flown. For application to Mars landers which must make use of the tenuous Martian atmosphere, it is critical to balance the lift of the vehicle to minimize the range error while still ensuring a safe deploy altitude.

Prior to the landing of Curiosity in August 2012, the Entry, Descent, and Landing (EDL) team examined minute tuning of the reference trajectory for the selected landing site in concert with other guidance parameters to balance the risk, and the vertical lift-to-drag ratio (L/D) command limits. The guided entry performed as designed without any significant exceptions. The initial guidance bank command was slightly larger than was expected due to the more dense atmosphere experienced prior to guidance start. Three bank reversals were executed and completed during the minute and a half long phase of range control. During the two minute duration of heading alignment, the entry vehicle steered directly towards the landing site. It delivered the rover to land 2.2 km away from the expected touchdown target. This miss distance is attributed to little time to correct the downrange drift from the final bank reversal and a suspected tailwind during heading alignment.

* Senior Engineer, Flight Dynamics Division, NASA Johnson Space Center, Houston TX 77058

† Aerospace Engineer, Flight Dynamics Division, NASA Johnson Space Center, Houston TX 77058

RELEVANT TERMINOLOGY

The following terms are common to entry guidance design and analysis and may be unfamiliar to some readers.

“In-plane” describes a vector component that is contained within the radius-velocity state vector plane using a planet-fixed coordinate system. This plane’s orientation changes slightly during entry as the vehicle’s azimuth varies. “Out-of-plane” describes a vector component that is normal to the same plane.

The term “vertical L/D” refers to the in-plane component of the L/D of the entry vehicle.

The term “downrange” describes the in-plane range from the entry vehicle state to the target position. “Crossrange” describes the out-of-plane range from the vehicle state to the target.

The term “bank angle” describes the rotation of the lift vector around the planet-relative velocity vector relative to the local horizon assuming a spherical planet. 0° is full lift-up, positive angles are to the right of the direction of flight. The lifting entry vehicle will vary the azimuth over time whenever it banks at values that are neither 0° nor 180° .

A “bank reversal” occurs when the sign of the commanded bank angle changes, indicating the bank direction of the vehicle should change from left to right of in-plane or vice versa.

“Planet-relative velocity” refers to the surface-relative velocity magnitude, using a planet-fixed coordinate system which accounts for planetary rotation. Any velocity reference in this paper is using this definition unless specifically defined as another. This velocity magnitude definition encompasses both the horizontal and vertical velocity components.

“Wind-relative velocity” refers to the airspeed of the entry vehicle, accounting for planetary rotation and the local wind velocity components.

ENTRY GUIDANCE OVERVIEW

MSL required a touchdown ellipse of 25 x 20 km for landing site selection. This led to a derived requirement to safely deploy the parachute within 10 km of the planned deploy target in order to achieve the required touchdown ellipse in the presence of winds. This deployment must occur in conditions that do not violate the parachute constraints and still allow sufficient time and altitude to complete the subsequent descent and landing tasks. Entry guidance must work in concert with the navigation and control systems to accomplish this.

Figure 1 describes the guided entry sequence of events. The Entry Terminal Point Controller (ETPC) entry guidance is divided into three distinct phases, discussed below in the order that they occur.

Pre-bank. The entry vehicle maneuvered into the pre-bank attitude minutes prior to entering the atmosphere. A pitch angle was commanded similar to the expected trim angle-of-attack when the atmosphere is encountered. The pre-bank angle was the bank angle maintained by the vehicle as it passed entry interface until range control began. If the predicted entry flight path angle was known to differ significantly from the reference trajectory, the pre-bank angle could be updated via remote command to minimize this initial bank maneuvering or control transients when range control started. This was intended to reduce the propellant usage by attempting to begin atmospheric flight near the trim angle of attack and the first commanded bank angle expected. The sign of the pre-bank angle, whether the vehicle banks left or right as it enters the atmosphere, was driven by communication link constraints with nearby Mars orbiters.

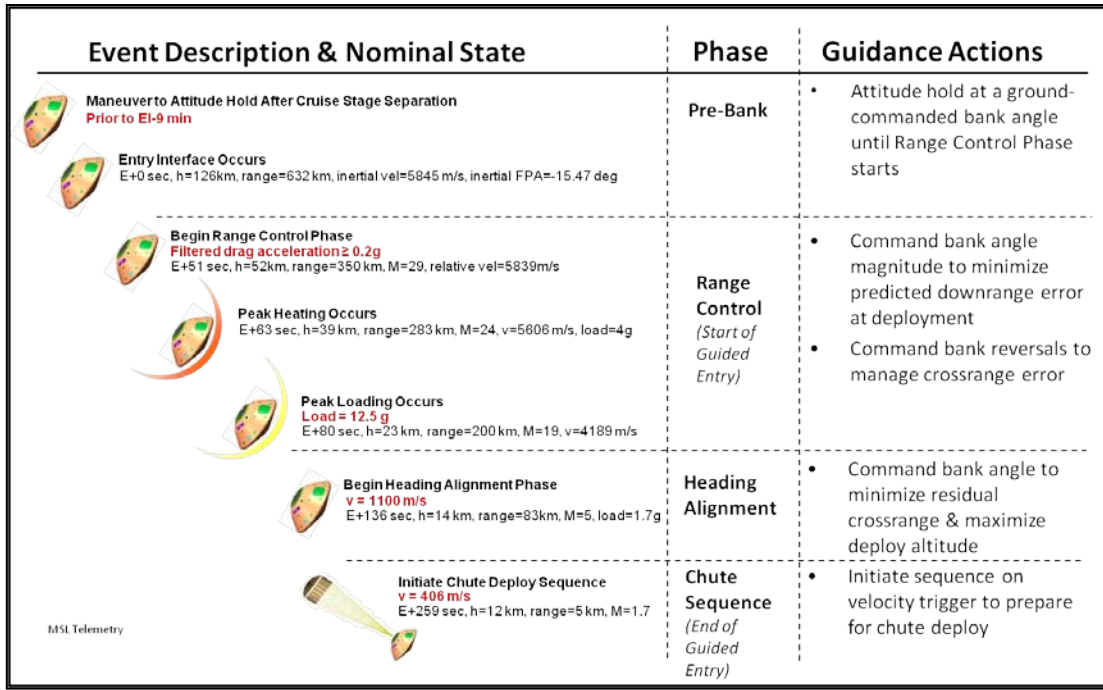


Figure 1. Sequence of entry guidance phases for an undispersed trajectory.

Range Control. Once the filtered drag acceleration magnitude climbed past 0.2 Earth g, the flight software determined that the vehicle has entered the sensible Martian atmosphere and began range control. During this phase the entry guidance predicted the downrange flown and commanded a bank angle to correct for any range errors. Simultaneously, the guidance monitored the crossrange to the target and commanded a bank reversal whenever the crossrange exceeded a deadband threshold. This ensured that the crossrange, although not directly controlled, was managed within a magnitude correctable during the next phase of heading alignment.

The range control logic has been detailed in previous papers^{1,2}.

The predicted range-to-go (R_p) was calculated as a function of filtered drag and altitude rate errors with the corresponding partial derivative gains with respect to the nominal reference trajectory profile, using Eq. (1). Experiencing less drag than the reference trajectory indicated the vehicle will fly farther than the reference. If the entry vehicle experienced a greater altitude rate than the reference, that indicated the vehicle will fly farther.

$$R_p = R_{ref} + \frac{\partial R}{\partial D} (D - D_{ref}) + \frac{\partial R}{\partial \dot{r}} (\dot{r} - \dot{r}_{ref}) \quad (1)$$

The desired vertical component of the lift-to-drag (L/D) ratio was calculated as a function of the difference between the actual (R) and predicted range-to-go (R_p), i.e., the downrange error. The difference was then converted to a change in vertical L/D commanded that is then applied to the reference vertical L/D at this velocity in Eq. (2). This included an overcontrol gain, K_3 , for robustness.

$$\left(\frac{L}{D}\right)_{V,C} = \left(\frac{L}{D}\right)_{V,ref} + \frac{K_3(R - R_p - R_{dep})}{\partial R / \partial (L/D)} \quad (2)$$

The commanded bank angle (Φ_C) was then calculated using the vertical L/D commanded and the estimated L/D as in Eq. (3). K_2 was -1 when banking left or 1 when banking right, and was changed upon a bank reversal command.

$$\Phi_C = \cos^{-1}\left(\frac{L/D_{V,C}}{L/D}\right) * K_2 \quad (3)$$

The bank reversals were commanded upon passing a quadratic threshold dependant on the velocity magnitude. The Apollo program found that a quadratic deadband, narrowing as the vehicle slowed, was a fair approximation of the remaining crossrange capability of the entry vehicle.

Heading Alignment. Once the estimated velocity dropped below 1100 m/s, the guidance ceased range control and began heading alignment. The commanded bank angle was now to steer the vehicle towards the target coordinates as seen in Eq. (4), where R_C is the target crossrange. Limiting the magnitude of the commanded bank angle to 30° ensured that most of the supersonic lift was countering gravity to prevent significant parachute deploy altitude loss.

$$\Phi_C = \tan^{-1}\left(\frac{R_c}{R}\right) K_4 \quad (4)$$

Guided entry ended when the sequence of events for parachute deploy was commanded, which included banking to 180° while jettisoning entry ballast to achieve a trim angle-of-attack near zero just prior to parachute deploy.

GUIDANCE PREPARATIONS FOR REAL-TIME OPERATIONS

The real-time operations during the week prior to landing required a substantial amount of preparation. During this week, the cruise and EDL teams were working together to determine whether any midcourse maneuvers were necessary prior to reaching Mars and which onboard parameters for navigation or guidance to update. The guidance parameters, including the pre-bank angle, the chute deploy range bias, and reference trajectory gains for range control were all candidates depending on the situation. The predicted entry flight path angle (EFPA) was the indicator used to determine the recommended action for any guidance parameter update.

Verification of ground commands played a role in the decision process. Changes to the pre-bank and deploy range bias were straight-forward and easy to error-check prior to being sent to the spacecraft. Changes to the reference trajectory gains required an update of several thousand parameters and required more effort to maintain the same diligence of error-checking. Consequently, the entry guidance team elected to pre-generate and certify a small set of alternate reference trajectory gain tables several months prior to real-time operations.

The subsequent section describes how the reference trajectory design map process¹ was modified to select reference trajectories for a particular entry flight path angle and resulting range to the landing target.

Navigation States

The interplanetary navigation team generated a number of trajectories from Earth departure to Mars entry in support of preparations for the guided entry. The reference design process required a parametric range of entry flight path angles (EFPA) to determine the reference trajectory and guidance performance with respect to incremental changes in EFPA. There is a strong correlation between EFPA and the initial range to the landing target; shallow entries are closer to the target when range control begins.

An intentional crossrange bias at atmospheric entry interface was implemented to ensure that the entry vehicle would avoid any descent stage debris during hypersonic flight. While a range of entry crossrange biases were analyzed in combination with the range of EFPA, it was determined that guidance performance was not sensitive to crossrange as long as it was within the crossrange capability of the vehicle.

Reference Design Process

The reference design process was repeated for a variety of combinations of different parameters describing the control trajectory, and resulted in numerous reference trajectories that were assessed using the entry simulations with embedded flight software. The process of parsing, filtering, and contour plotting of the data to determine the best reference trajectory was referred to as the reference design map¹. The open-loop shaping of the profile is examined with a parametric study over a range of bank angles. A simple variable bank profile is that of a linear ramp between two constant values: an early bank/velocity near the start of guidance range control and a late bank/velocity near the end of range control². For the construction of the guidance reference design map, each EFPA and reference bank profile was parametrically swept through a range of values. While an optimizer could be employed, the parametric approach is sufficient and provides insight as to the behavior across the full range.

Each reference trajectory was used to generate a set of influence coefficients commonly referred to as the guidance gains². The closed loop performance in stressing trajectory conditions was assessed for each reference trajectory of the parametric sweep¹. Performance parameters during the guided entry, such as the parachute deploy altitude and miss distance, were filtered to identify cases meeting minimum constraints. Those cases were examined to choose the best bank profile for guidance performance across a particular EFPA range.

Since parachute deploy altitude was historically a driving constraint in EDL performance, selecting the best reference bank profile was based in large part upon maximizing altitude while minimizing guidance errors. An additional filter based on the range flown from entry to the target was used for real-time operations because now it is also important to have a similar range flown so that the guidance gains are reliable. The dots in Figure 3 illustrate how the range flown in the various reference trajectories is a function of the EFPA, and should be compared with the heavy dashed lines that represent the actual range to the target given a particular EFPA. If the EFPA is different from the nominal value of -15.5° , then the additional constraint of the range flown results in reference trajectories that are no longer as optimal for altitude but will still land inside the target ellipse.

Any trajectory case that did not meet the filters was rejected out of hand. Of those remaining, the case with the highest chute deploy altitude was identified and its bank profile was noted. This case was plotted and compared against previous reference trajectories, as well as checked for any desirables or discrepancies in guidance parameters and performance.

Altitude-optimized bank profiles were found for each EFPA. Each reference trajectory from a particular EFPA was run in a Monte Carlo simulation with different EFPA to examine the performance and identify which gain sets best worked across the range of EFPA. Figure 4 shows how guidance performance and delivery error are affected based upon using a baseline and updated reference gain set. To simplify operations, it was decided that only three reference gain sets were needed to span the range of feasible EFPAs during nominal and off-nominal operations. For each gain set, a table was printed out for real-time operations that provided the best pre-bank angle and chute deploy range bias as a function of EFPA. These tables were certified for the EDL team to use during operations. If the situation did not permit the uplink of an updated reference gain set, the tables included contingency adjustments of the chute deploy range bias and pre-bank angle to recover some guidance performance and allow a high probability of mission success.

GUIDANCE FLIGHT OPERATIONS




Figure 5. Entry Point Uncertainties from E-8 days to E-1 day within EFPA corridor.

During the last week of operations prior to landing, there was continuous tracking of the spacecraft as it approached Mars in order to provide the best prediction of its health and atmospheric entry point. For each navigation orbital determination (OD) assessment, a 3-sigma uncertainty ellipse was estimated. Figure 5 shows these OD solutions charted in a B-plane to

ascertain the entry point location with respect to the EFPA corridor. The B-plane is useful to describe the approaching trajectory of the spacecraft relative to the planet³. The yellow and green lines each show the corridor in which TCM-5 is acceptable (yellow) and optimal (green). Each successive OD solution from the first onboard parameter update through just prior to entry shows the uncertainty in the entry state shrinking as knowledge improved. OD entry point solutions shifting to the right in the plot are shallower than the nominal EFPA, and solutions shifting to the left are steeper than the nominal EFPA. This plot was utilized as a decision-making guide whether to perform a midcourse maneuver (TCM), or update the onboard parameters (among them entry guidance parameters). If any of the OD solutions trended outside the aforementioned corridors, the guidance parameters would be updated enacted from the lookup tables³.

In addition to examining EFPA, nominal and operations Monte Carlo dispersed trajectories were run to verify that EDL performance was as expected.

The final MSL approach eight days before entry had OD solutions yielding an EFPA error of about 0.16° . It was decided that TCM-4 would be performed. After the successful execution of this maneuver and OD solutions thereafter, the first onboard parameter update was executed³. This parameter update was mandatory since no onboard nav state for EDL existed onboard the spacecraft. In order to optimize guidance performance for this spacecraft state, the prebank and ranoff values were uplinked as part of the onboard parameter update five days before entry. Post uplink, it was verified that the nav and guidance parameters were successfully input to the flight software onboard. Simulations showed consistency in expected guidance performance. After this first update, OD solutions continued to converge and the entry point remained on-target. Though assessment continued, because of the high degree of accuracy and stability of those solutions, it was decided that no further TCMs or onboard parameter updates were necessary before landing.

RECONSTRUCTED ENTRY TRAJECTORY AND FLIGHT RESULTS

The reconstructed entry trajectory, from entry interface down to the start of the parachute deploy sequence, is shown in Figure 6 along with annotations of when pre-bank, range control, and heading alignment phases occurred. The time is relative to when the EDL flight software initialized with the onboard state vector previously provided from the ground. This initialization occurred approximately nine minutes prior to atmospheric entry. The entire duration of the guided entry was over three and a half minutes, with less than half of that time spent in range control. Geodetic altitude is referenced to the Mars MOLA areoid surface definition. The landing site at Gale Crater has an elevation near -4.6 km relative to that areoid.

The estimated bank plot at the lower right of Figure 6 clearly depicts the three bank reversals that occur during range control. During heading alignment the bank angle briefly lingered near -30° and then steadied near 0° while guidance finished the alignment towards the landing site. The minute changes in actual bank after that are primarily due to the estimated bank angle reaching the attitude controller deadbands. The sudden shift towards -180° bank after 780 seconds denotes the start of the re-orientation just prior to parachute deploy.

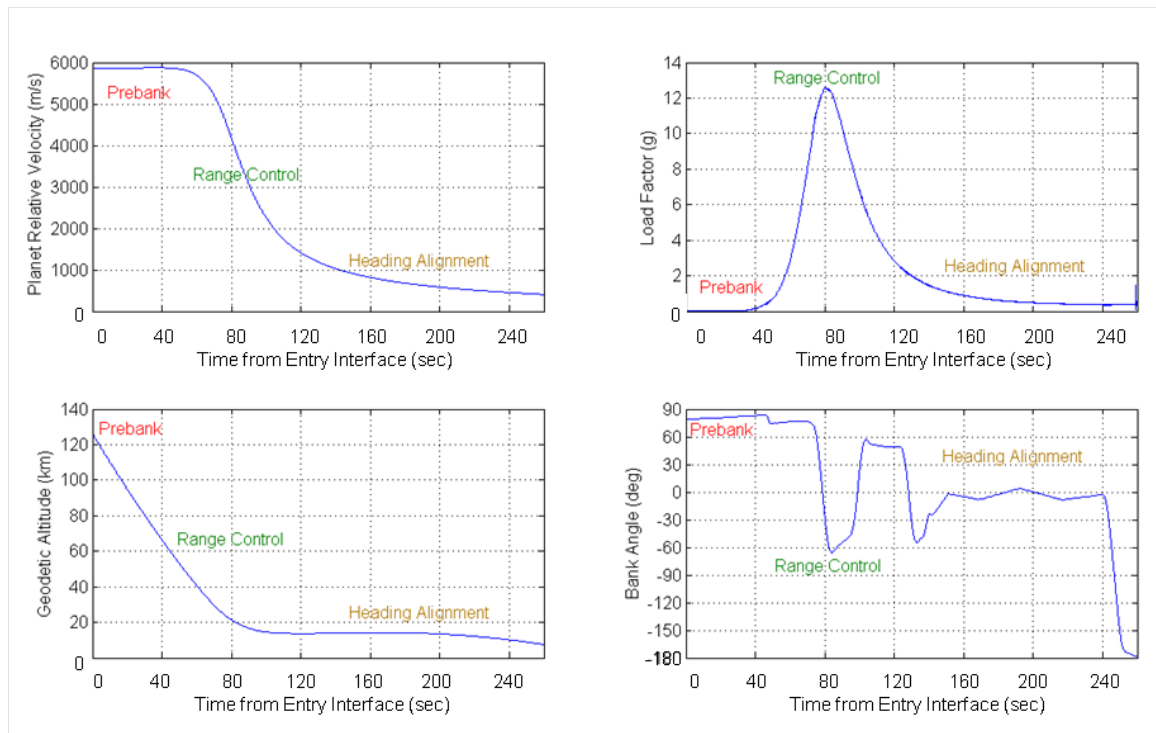


Figure 6. Best estimate of the entry trajectory.

A sample of entry guidance performance plots of the flight data are given in the next set of figures. The subsequent figures represent flight telemetry data plotted versus estimated velocity and time relative to EDL state initialization, with both axes progressing in time from left to right to ease the user in correlating events between the two. Parameters relative to banking or range error have been oriented so that that lift-up values are towards the top of the plot. These unconventional axes have proven to be quite readable among the interdisciplinary EDL team. Gray areas represent the duration of bank reversals and other slews, where the vehicle was largely unresponsive to guidance commanding.

One of the first plots examined from the entry guidance flight software telemetry is the commanded bank profile during range control, depicted in Figure 7. To provide sufficient detail on the commanded bank angle variations, the magnitude is plotted and the color coding of the line denotes whether the guidance is commanded a negative bank to the left of the plane of motion or a positive bank to the right. The dashed line is the vertical L/D from the reference trajectory, converted to bank angle using the onboard, filtered L/D at that time. When the commanded bank magnitude is near the reference line there is little predicted range error the guidance is attempting to correct.

The first bank command was roughly 10° away from the pre-bank command. It is not apparent from this plot, but this difference is attributed to the greater atmospheric density experienced at high altitudes⁴. During the bank reversals the guidance continuously updated the commanded bank angle for the attitude controller. Since the range control algorithm includes altitude rate errors, very quickly the guidance began to respond to the altitude rate perturbation introduced by each bank reversal. It is apparent that the guidance is converging back towards the reference trajectory, suggesting range convergence due to the overcontrol gain. While the commanded bank profile is useful, in itself it doesn't inform what the range control algorithm is responding to. This deficiency is pronounced in dispersed cases or flight data and makes it impractical to diagnose the guidance response from commanded bank angle alone.




Figure 7. Commanded bank angle during range control.

A predicted range error component plot, Figure 8, illustrates which errors the guidance output is responding to and provides a more thorough explanation of the factors at work. The predicted range errors are summed as shown below in Eq. (5). The downtrack error component is the difference of the current in-plane range to the deploy target versus the range-to-go that the reference trajectory table has at this velocity and the deploy range bias. The drag error component is the difference between the current filtered drag versus the reference drag at this velocity, with the reference drag-to-range partial applied. The r-dot error component is the altitude rate difference compared to the reference at this velocity, with the reference r-dot-to-range partial applied. The summation of all three components provides the total predicted range error that the guidance is using in Eq. (2) to determine the vertical L/D command. The area plot shows the summation of these components, including how one may cancel another out, in the order shown in the legend. Positive values indicate the vehicle will land uprange of the target.

$$R_{error} = \left(R - R_{ref} - R_{dep} \right) - \frac{\partial R}{\partial D} \left(D - D_{ref} \right) - \frac{\partial R}{\partial \dot{r}} \left(\dot{r} - \dot{r}_{ref} \right) \quad (5)$$

It is useful to walkthrough one instance so to ensure proper interpretation of this plot. At 5500 m/s, the downtrack error component is near -5 km. The drag predicted range error component is an additional -3 km. The r-dot predicted range error component is +2 km. This results in a net of -6 km of predicted range error that the guidance is correcting at that instant. The effect of the first bank reversal on the range error is obvious just as it was in the commanded bank plot. However, now we can clearly observe that the guidance, due to the range overcontrol gain, is quickly correcting the reversal perturbation and driving the predicted range error back towards zero before the next reversal occurs.

By the end of the range control phase, the downtrack error component is near -3 km, the drag error component is near +1 km and the r-dot error component is nearly 0. This results in a final predicted range error near -2 km. It is clear that the final bank reversal introduced some additional predicted range error, as expected, and that there was not enough time left during the range control phase to significantly reduce that impact. By the end of range control, there was a predicted miss of over 1 km past the target.

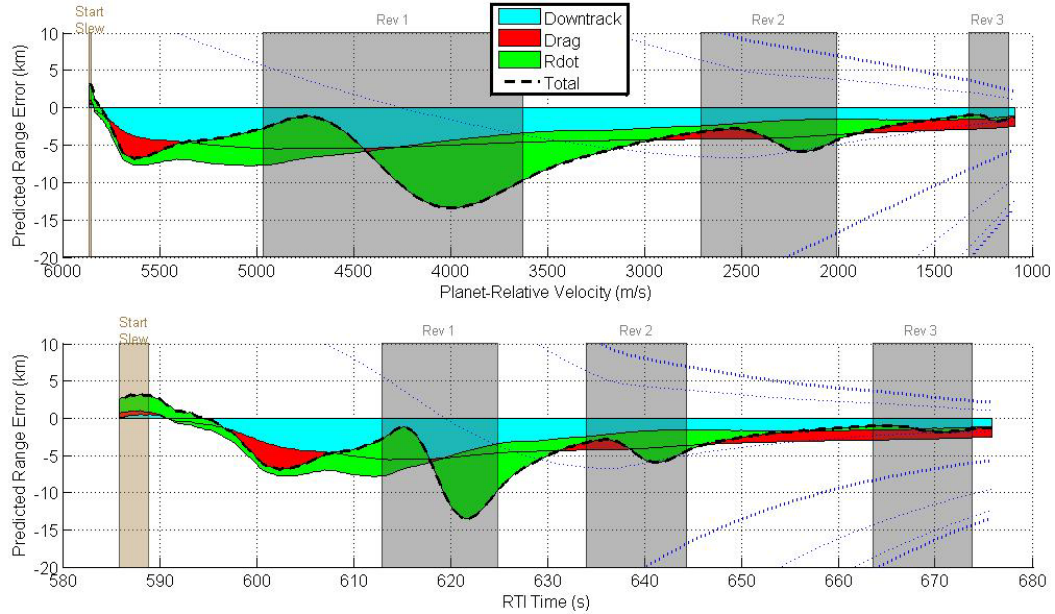


Figure 8. Range error components observed during range control.

Figure 9 illustrates the estimated lateral angle, representing crossrange, as a function of time relative to the lateral deadbands that initiate bank reversals. The lateral deadband for the first reversal is tighter and then relaxes to a wider deadband for the remaining reversals^{1,2}. It is clear that the bank reversals start when the lateral angle estimate crosses the deadband.




Figure 9. Bank reversal lateral angle during range control segment

The commanded bank angle during heading alignment in Figure 10 is fairly simple compared to range control. For a few seconds in at the beginning, the guidance commanded as much left bank as allowed until the estimated location of the target site was approaching in-plane. After that, the commanded bank angle slowly oscillates near zero depending on where the actual bank angle was relative to the entry controller deadband. The bias of this average bank angle may be due to some combination of a crosswind component from the north (left of the plane of motion) or a sideslip angle^{4,5}.




Figure 10. Commanded bank angle observed during heading alignment segment

Understanding of the impact of various error sources on the actual deploy and landing targets is achieved by incrementally replacing the dispersed models in the Monte Carlo with reconstructed models to chart the displacement of the mean⁶. Table 1 lists each data source with its associated predicted error from guidance and mean downrange at chute deploy and landing as different models are replaced in the pre-entry Monte Carlo with reconstructed models. Positive downrange values indicate that point was uprange or landing short. These values are normalized relative to the pre-entry Monte Carlo to demonstrate the effect compared to the EDL team predictions. The delta captures the difference from that reconstructed model to the line above. The predicted guidance downrange error occurs at the end of the range control phase, two minutes before parachute deploy.

Navigation uncertainties, combining both delivery and onboard knowledge error, contribute about 0.1 km of the total error at parachute deploy and touchdown. Entry guidance cannot correct for onboard knowledge, as guidance utilizes this knowledge. Reconstructed aerodynamic properties⁴ appear to push deploy and landing points uprange of the target approximately 0.7 km from the reconstructed navigation. Additional uncertainties not yet modeled, including atmosphere and winds, contribute 1 km of error from the aero target to the as-flown chute deploy target. The potential impact of atmospheric and wind impact on downrange were expected¹ and it is unsurprising to see the magnitude of this effect⁵. Further studies are in work to reconstruct the atmosphere and winds to better characterize the impact on parachute deploy and landing.

The reported guidance accuracy is the reconstructed downrange after subtracting out the navigation errors which guidance cannot correct. The rover landed 2.3 km downrange of the pre-entry target that the mission was aiming at, but only 2.2 km downrange from the expected touchdown target when accounting for the onboard knowledge error.

Table 1: Cumulative Sources of Error on Landing Accuracy

Source Data	Guidance Prediction, Relative (km)	Delta (km)	Downrange at Deploy, Relative (km)	Delta (km)	Downrange at Landing, Relative (km)	Delta (km)
Monte Carlo of Pre Entry Trajectory	0.0	--	0.0	--	0.0	--
+ Actual EFPA Delivery	0.0	0.0	0.0	0.0	0.0	0.0
+ Actual Onboard Nav State	-0.1	-0.1	-0.1	-0.1	-0.1	-0.2
+ Reconstructed Aerodynamics	0.5	0.5	0.7	0.8	0.6	0.7
Best Estimate from Reconstructed Trajectory	-0.6	-1.0	-2.4	-3.1	-2.3	-2.9

Figure 11 charts the various source deploy targets in areocentric latitude and longitude space. The spacecraft approach vector is from the west-north-west parallel to the ellipse major axis. Examination of the flight telemetry did find an uncharacteristically long time period between the start of the re-orient sequence and parachute deploy. This suggests low drag was experienced as both events were triggered based off of estimated velocity. There is some question as to the accuracy of the supersonic aerodynamic reconstruction which increased the capsule drag coefficient compared to the undispersed prediction⁴. If the coefficient reconstruction is correct, the dynamic pressure must have been lower. This, along with the reconstructed deploy point, point to a tailwind during heading alignment and the re-orient sequence.




Figure 11. Contributions to chute deploy target error




Figure 12. Contributions to landing target error

The actual chute deploy occurred about 2.4 km away from the pre-entry target, and additional dispersions pushed the landing target 2.3 km, as seen in Figure 12. It is suggested that a tailwind could have been imparted prior to and following chute deploy that led to the actual landing point being further downrange than predicted.

CONCLUSION

The first guided entry at Mars was a success, landing Curiosity safely inside the landing ellipse only 2.2 km away from the expected point and well within the one-sigma ellipse and landing ellipse requirements. A late bank reversal and a suspected tailwind near the end of guided entry contributed to this slight miss. Further characterization of the atmospheric and wind effects on the landing position and will aid in the understanding of the observed entry guidance response. Refinement of the guidance gains and alternative parachute deploy sequence triggers to reduce the landing ellipse size will be studied for future Mars landing missions.

ACKNOWLEDGEMENTS

Mars Science Laboratory is managed by the Jet Propulsion Laboratory under contract with NASA. NASA Johnson Space Center is tasked by JPL to provide entry guidance expertise in design, analysis, and operations for MSL. The authors would like to acknowledge their colleagues on the MSL Entry, Descent, and Landing Team whom made significant contributions to this effort: A.M.S. Martin, A. Chen, and P. Brugarolas of JPL, and D.W. Way, R.W. Powell, J.D. Shidner, J.L. Davis, and A.M. Dwyer-Cianciolo of NASA Langley Research Center.

The participation by NASA Johnson Space Center on this successful mission would not have been possible without Claude Graves and the exemplary technical work by Gilbert Carman prior to his retirement in 2005.

REFERENCES

- ¹ Mendeck, G.; Craig, L., "Entry Guidance for the 2011 Mars Science Laboratory Mission," AIAA 2011-6639, AIAA Atmospheric Flight Mechanics Conference, August 2011, Portland, OR.
- ² Mendeck, G.; Carman, G., "Guidance Design for Mars Smart Landers Using the Entry Terminal Point Controller," AIAA 2002-4502, AIAA Flight Mechanics Conference, August 2002, Monterey, CA.
- ³ Martin-Mur, T.; Kruizinga, G. "Mars Science Laboratory Interplanetary Navigation Performance," AAS 13-232, AIAA Spaceflight Mechanics Conference, February 2013, Lihue, Kauai.
- ⁴ Schoenenberger, M., Van Norman, J. "Preliminary Trajectory Reconstruction Results of the Mars Science Laboratory Entry Vehicle," AAS 13-306, AIAA Spaceflight Mechanics Conference, February 2013, Lihue, Kauai.
- ⁵ Karlgaard, C.; Schoenenberger, M. "Mars Science Laboratory Entry, Descent, and Landing Trajectory and Atmosphere Reconstruction," AAS 13-307, AIAA Spaceflight Mechanics Conference, February 2013, Lihue, Kauai.
- ⁶ Euler, E.A.; Adams, G.L. "Design and Reconstruction of the Viking Lander Descent Trajectories," AIAA 55795-480, AIAA/AAS Astrodynamics Specialist Conference, September 1977, Jackson Hole, WY.